

WHAT IS CLAIMED IS:

1. A system comprising:
a network;
5 a plurality of computing nodes coupled via the network;
wherein a first node from the plurality of nodes is operable to pre-allocate a plurality of data object replicas;
wherein the first node is operable to use the plurality of data object replicas to satisfy a request to create a first data object in response to receiving the request after said
10 pre-allocating the replicas.
2. The system of claim 1,
wherein said using the plurality of data object replicas to satisfy the request to create the first data object comprises associating the plurality of data object replicas with
15 the first data object.
3. The system of claim 1,
wherein for each data object replica, said pre-allocating the data object replica comprises allocating a portion of disk space for the data object replica.
20
4. The system of claim 1,
wherein said pre-allocating the plurality of data object replicas comprises pre-allocating a data object replica on each of a plurality of nodes.
- 25 5. The system of claim 4,
wherein the first node is further operable to select the plurality of nodes on which to pre-allocate the data object replicas.
6. The system of claim 5,

wherein the first node maintains first information indicative of free space on other nodes;

wherein said selecting the plurality of nodes on which to pre-allocate the data object replicas comprises selecting the plurality of nodes based on the first information.

5

7. The system of claim 1,

wherein said pre-allocating the data object replicas decreases latency of the request to create the first data object by enabling the first node to satisfy the request without performing replica allocation in response to the request.

10

8. The system of claim 1,

wherein said pre-allocating the data object replicas comprises storing information enabling nodes to access the replicas.

15

9. The system of claim 1,

wherein the first node is operable to maintain a cache of pre-allocated replicas, wherein the cache comprises multiple sets of pre-allocated replicas.

10. The system of claim 1,

20

wherein the request to create the first data object specifies a name for the first data object;

wherein said satisfying the request to create the first data object comprises associating the name with the plurality of data object replicas.

11. The system of claim 1,

25

wherein after said using the plurality of data object replicas to satisfy the request to create the first data object, the first node is operable to pre-allocate another plurality of data object replicas to make the other plurality of data object replicas available for a subsequent request to create another data object.

30

12. A system comprising:
a network;
a plurality of computing nodes coupled via the network;
5 wherein a first node from the plurality of nodes is operable to pre-allocate a plurality of data object replicas for use in satisfying a subsequent request to create a new data object;
wherein said pre-allocating the plurality of data object replicas comprises:
selecting a first plurality of nodes on which to allocate the replicas;
10 and
allocating a replica on each of the first plurality of nodes.

13. The system of claim 12, wherein said pre-allocating the plurality of data object replicas further comprises:
15 generating a first data object ID; and
associating the first data object ID with each of the plurality of data object replicas.

14. The system of claim 12, wherein said pre-allocating the plurality of data object replicas further comprises:
20 storing information enabling nodes to access the plurality of data object replicas.

15. The system of claim 12, wherein said pre-allocating the plurality of data object replicas further comprises:
25 storing information to link the plurality of data object replicas together.

16. A carrier medium comprising program instructions executable to
30 implement the method of:

a first node from a plurality of nodes pre-allocating a plurality of data object replicas;

the first node receiving a request to create a first data object after said pre-allocating the replicas; and

5 the first node satisfying the request to create the first data object using the plurality of data object replicas.

17. The carrier medium of claim 16,

wherein said satisfying the request to create the first data object using the
10 plurality of data object replicas comprises associating the plurality of data object replicas with the first data object.

18. The carrier medium of claim 16,

wherein for each data object replica, said pre-allocating the data object replica
15 comprises allocating a portion of disk space for the data object replica.

19. The carrier medium of claim 16,

wherein said pre-allocating the plurality of data object replicas comprises pre-allocating a data object replica on each of a plurality of nodes.

20

20. The carrier medium of claim 19,

wherein the first node also selects the plurality of nodes on which to pre-allocate the data object replicas.

25 21. The carrier medium of claim 20,

wherein the first node also maintains first information indicative of free space on other nodes;

wherein said selecting the plurality of nodes on which to pre-allocate the data object replicas comprises selecting the plurality of nodes based on the first information.

30

22. The carrier medium of claim 16,
wherein said pre-allocating the data object replicas decreases latency of the request to create the first data object by enabling the first node to satisfy the request without performing replica allocation in response to the request.

5

23. The carrier medium of claim 16,
wherein said pre-allocating the data object replicas comprises storing information enabling nodes to access the replicas.

10

24. The carrier medium of claim 16,
wherein the first node maintains a cache of pre-allocated replicas, wherein the cache comprises multiple sets of pre-allocated replicas.

15

25. The carrier medium of claim 16,
wherein the request to create the first data object specifies a name for the first data object;

wherein said satisfying the request to create the first data object comprises associating the name with the plurality of data object replicas.

20

26. The carrier medium of claim 16,
wherein after said satisfying the request to create the first data object using the plurality of data object replicas, the first node pre-allocates another plurality of data object replicas to make the other plurality of data object replicas available for a subsequent request to create another data object.

25

27. A carrier medium comprising program instructions executable to implement the method of:

pre-allocating a plurality of data object replicas for use in satisfying a subsequent request to create a new data object;

30

wherein said pre-allocating the plurality of data object replicas comprises:
selecting a first plurality of nodes on which to allocate the replicas;
and
allocating a replica on each of the first plurality of nodes.

5

28. The carrier medium of claim 27, wherein said pre-allocating the plurality of data object replicas further comprises:

generating a first data object ID; and
associating the first data object ID with each of the plurality of data object
10 replicas.

29. The carrier medium of claim 27, wherein said pre-allocating the plurality of data object replicas further comprises:

storing information enabling nodes to access the plurality of data object
15 replicas.

30. The carrier medium of claim 27, wherein said pre-allocating the plurality of data object replicas further comprises:

storing information to link the plurality of data object replicas together.

20